



(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
16.05.2018 Bulletin 2018/20

(51) Int Cl.:
G10L 17/24 (2013.01) **G06F 21/46** (2013.01)
G10L 15/08 (2006.01) **G10L 15/22** (2006.01)
G10L 25/51 (2013.01) **G06F 21/32** (2013.01)

(21) Application number: **17198918.9**

(22) Date of filing: **01.07.2013**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

• **SCHALKWYK, Johan**
Scarsdale, NY New York 10583 (US)
• **SAN MARTIN, Maria Carolina Parada**
Palo Alto, CA California 94303 (US)

(30) Priority: **03.07.2012 US 201261667817 P**
06.08.2012 US 201213567572

(74) Representative: **Chettle, John Edward et al**
Venner Shipley LLP
200 Aldersgate
London EC1A 4HD (GB)

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
13739887.1 / 2 870 602

(71) Applicant: **Google LLC**
Mountain View, CA 94043 (US)

Remarks:

This application was filed on 27-10-2017 as a divisional application to the application mentioned under INID code 62.

(72) Inventors:
• **RUBIN, Andrew E.**
Portola Valley, CA California 94028 (US)

(54) **DETERMINING HOTWORD SUITABILITY**

(57) Methods, systems, and apparatus, including computer programs encoded on a computer storage medium, for determining hotword suitability. In one aspect, a method includes receiving speech data that encodes a candidate hotword spoken by a user, evaluating the speech data or a transcription of the candidate hotword,

using one or more predetermined criteria, generating a hotword suitability score for the candidate hotword based on evaluating the speech data or a transcription of the candidate hotword, using one or more predetermined criteria, and providing a representation of the hotword suitability score for display to the user.




FIG. 1

Description

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims the benefit of U.S. Provisional Pat. Application No.: 61/667,817, filed July 3, 2012, and U.S. Pat. Application No. 13/567,572, filed August 6, 2012, the entirety of both applications are hereby incorporated by reference in their entirety.

BACKGROUND

[0002] The present specification relates to the determination of hotword suitability.

[0003] Hotwords, sometimes referred to as "attention words" or "wake-up words," provide a hands-free approach for a user to begin an interaction with a speech-enabled computing device. When the computing device detects that the user has spoken a hotword, the device enters a ready state for receiving further user commands. Device manufacturers and software creators generally provide predetermined hotwords to be employed by users for interacting with their devices or software applications.

SUMMARY

[0004] According to one innovative aspect of the subject matter described in this specification, a user of a mobile device may define their own hotword to be used by the mobile device. A candidate hotword may be evaluated to determine, for example, whether the candidate hotword may be easily confused by the device with another word, or whether the candidate hotword may be inappropriate for use as a hotword. A result of this evaluation is displayed to the user, who may then confirm whether they would indeed like to designate the candidate hotword as the hotword, or submit a different candidate hotword for evaluation.

[0005] As another innovative aspect of the subject matter described in this specification, a process includes receiving speech data that encodes a candidate hotword spoken by a user, evaluating the speech data or a transcription of the candidate hotword using one or more predetermined criteria, generating a hotword suitability score for the candidate hotword based on evaluating the speech data or a transcription of the candidate hotword using one or more predetermined criteria, and providing a representation of the hotword suitability score for display to the user.

[0006] Other embodiments of these aspects include corresponding systems, apparatus, and computer programs, configured to perform the actions of the methods, encoded on computer storage devices.

[0007] These and other embodiments may each optionally include one or more of the following features. For instance, the hotword suitability score reflects a suitability of the candidate hotword for use as a hotword; a hotword

is a keyword or a key phrase for waking up a device; a hotword is a keyword or a key phrase for unlocking a device; a hotword is a voice password; the speech data is received from a mobile device that is operated by the user; generating a hotword suitability score includes generating, for each of two or more predetermined criteria, a feature score based on an evaluation of the speech data or of the transcription of the candidate hotword using the predetermined criteria, aggregating the feature scores, and outputting, as the hotword suitability score, the aggregated feature scores; evaluating the speech data or a transcription of the candidate hotword includes determining a word length, in frames, of the speech data, where longer word lengths are associated with higher hotword suitability scores; evaluating the speech data or a transcription of the candidate hotword includes determining a quantity of syllables or phones associated with the candidate hotword, where a larger quantity of syllables or phones is associated with a higher hotword suitability score; evaluating the speech data or a transcription of the candidate hotword includes determining an edit distance between the transcription of the candidate hotword and one or more other words that share a same language as the candidate hotword, where smaller edit distances are associated with lower hotword suitability scores; evaluating the speech data or a transcription of the candidate hotword includes determining a quantity of pronunciations specified for the candidate hotword, where fewer quantities of pronunciations are associated with higher hotword suitability scores; evaluating the speech data or a transcription of the candidate hotword includes determining that one or more particular phones occur in the transcription of the candidate hotword; the occurrence of the particular phones in the transcription is associated with a higher hotword suitability score; the occurrence of the particular phones in the transcription is associated with a lower hotword suitability score; evaluating the speech data or a transcription of the candidate hotword includes determining a quantity of training examples of the hotword that are available to a speech recognizer, where a smaller quantity of training examples is associated with a lower hotword suitability score; evaluating the speech data or a transcription of the candidate hotword includes determining an extent to which pronunciation of the candidate hotword varies among multiple utterances of the candidate hotword by the user, where a greater extent is associated with a lower hotword suitability score; the representation of the hotword suitability score includes data indicating that the candidate hotword is suitable as a hotword; the representation of the hotword suitability score includes data indicating that the candidate hotword has been accepted as a hotword; the representation of the hotword suitability score includes a color-coded representation of the hotword suitability score; and/or the representation of the hotword suitability score includes a bar, a gauge or a meter.

[0008] Advantageous implementations may include one or more of the following features. Criteria related to

a candidate hotword and to a particular speaker (i.e., a device user) may be considered when evaluating the hotword. Various hotword suitability criteria may be scored and aggregated, ensuring that each criterion is considered, and that a particular criterion does not dominate during a hotword evaluation process. Device users may receive hotword suitability scores for making informed decisions when selecting hotwords for their devices. The use of suitable hotwords may improve the ability of a device to perform hotword detection, thus improving device performance and user experience. Moreover, the use of suitable hotwords may limit the potential for a "false positive" hotword detection to mistakenly activate a device, thus conserving device resources.

[0009] The details of one or more implementations are set forth in the accompanying drawings and the description below. Other potential features and advantages will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010]

FIGS. 1 and 2 are diagrams of example systems that may determine the suitability of a spoken hotword.

FIG. 3 is a flow chart illustrating an example process for determining the suitability of a spoken hotword and for providing a representation of the suitability.

FIG. 4 shows an example of a generic computer device and a generic mobile computer device.

[0011] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0012] FIG. 1 is a diagram of an example system 100 that may determine the suitability of a spoken hotword. FIG. 1 also illustrates an example flow of data, shown in stages (A) to (H). Stages (A) to (H) may occur in the illustrated sequence, or they may occur in a sequence that is different than in the illustrated sequence.

[0013] In general, hotword suitability may be determined by the system 100 as part of a hotword designation process. Hotwords may include keywords or key phrases spoken by users for interacting with or controlling stationary, portable, or mobile computing devices. For example, a hotword may be spoken by a user to initiate communication with a computing device, generally placing the device in a state in which it may receive further voice commands. Upon detecting a spoken hotword, for example, a computing device may perform various actions, including activation (i.e., "waking up" from a dormant or idle state), unlocking, launching a device application (e.g., a search application, a map application, etc.), loading pre-

ferred user settings, and/or other appropriate actions.

[0014] Device users may want to select particular words or phrases to serve as hotwords for their computing devices. For example, a personally meaningful word or phrase may serve to customize a device, may personalize the device in a humorous way, and may be easier for a user to remember or pronounce than a word or phrase provided by a device manufacturer. Moreover, user-selected hotwords may provide a measure of security, essentially functioning as a password. Some words or phrases, however, may be more readily detected by a computing device than others, and thus may serve as better hotwords. As will be described in examples below, hotword detectability may depend on various criteria related to a particular word or phrase and related to a particular speaker, such as the length of the word or phrase, the quantity of syllables or phones included in the word or phrase, the occurrence of particular phones, the quantity of identified pronunciations for the word or phrase, the variability of word or phrase pronunciation by the speaker, the availability of training examples for the word or phrase from various other speakers, and the edit distance between a spoken word or phrase and other words or phrases of a shared language. Each of the criteria may be scored, aggregated, and output to a user, thus providing an overall indication of the suitability of a candidate hotword. Based on hotword suitability scores, for example, users may make informed decisions when selecting hotwords for their devices.

[0015] In further detail, the system 100 includes a computing device 102 (shown here in various stages of operation as computing devices 102a, 102b, and 102c) which can present an interface 104 (shown here in various states of operation as interfaces 104a, 104b, and 104c) for display to a user. The computing device 102 includes one or more processors configured to execute instructions stored by a computer-readable medium for performing various device operations, such as input/output, communication, data processing, and the like. For example, the computing device 102 may include one or more client computing devices (e.g., desktop computers, laptop computers, notebook computers, navigation systems, smartphones, or other appropriate computing devices) and/or may include one or more servers (e.g., a backend server farm). Input received by a client computing device, for example, may be processed locally, or may be provided to a backend system (e.g., the cloud) for processing. As another example, processing may be distributed between one or more client computing devices (e.g., in a peer-to-peer configuration) and/or servers.

[0016] Referring to the example data flow, during stage (A), a user is prompted to speak or type a candidate hotword. For example, the computing device 102 (shown here as device 102a) may provide an instructional message to the user via the interface 104 (shown here as interface 102a) to "Please say your desired hotword". In the present example, the user utters the candidate hotword "pizza" and the computing device 102 captures,

encodes, and/or generates speech data 110 based on the utterance.

[0017] The speech data 110, for example, may include audio source data 112 (e.g., an audio stream, a digitally sampled audio file, or another appropriate data format), a corresponding set of phones 114 (i.e., speech sounds), and optionally, a transcription 116. To identify the set of phones 114 (e.g., speech sounds detected in the user's utterance of "pizza"), for example, the computing device 102 and/or backend computing devices can analyze acoustical signals included in the audio source data 112, and can compute and classify sound features based on the analysis. Further, to generate the transcription 116 (e.g., the transcribed word "pizza"), the computing device 102 and/or backend computing devices can use speech recognition techniques (e.g., Hidden Markov models, neural networks, or other appropriate techniques) to associate the set of phones 114 with a word or phrase. Upon capturing, encoding, and/or generating the speech data 110, for example, the data may be provided to one or more hotword strength evaluation engines 120.

[0018] During stage (B), speech data is processed and evaluated. In the present example, the hotword strength evaluation engines 120 may process and evaluate the speech data 110. The hotword strength evaluation engines 120, for example, may include one or more hardware and/or software components (e.g., modules, objects, libraries, etc.) configured to process and evaluate speech data, and may be executed by one or more processors of the computing device 102 and/or other associated devices (e.g., backend servers, peer devices, etc.).

[0019] In general, various criteria may be used for evaluating speech data. Each of the criteria, for example, may be associated with a different hotword strength evaluation engine, which may in turn produce a feature score based on a particular criterion. Each of the different hotword strength evaluation engines, for example, may evaluate audio source data, derived sets of phones, transcriptions, and/or other appropriate speech data associated with a candidate hotword. Moreover, in some implementations, training examples (speech data associated with past utterances of the candidate hotword by the speaker and/or by other speakers) may be evaluated. Training examples, for example, may be maintained by a training examples data store 122. The training examples data store 122 can implement databases, data services, file systems, and the like to add, remove, and maintain data used by the system 100.

[0020] In the present example, one or more of the hotword strength evaluation engines 120 may evaluate the audio source data 112 to determine a word length in frames (i.e., a time-based section of audio). For example, the candidate hotword "pizza" may be considered to be relatively short compared to other words, or the particular speaker may happen to utter the word "pizza" relatively quickly compared to other words or other speakers. Based on a word length criterion, for example, one of the

hotword strength evaluation engines 120 may produce a corresponding feature score 124a (e.g., a low score).

[0021] Continuing with the present example, one of the hotword strength evaluation engines 120 may evaluate the set of phones 114 and/or the transcription 116 to identify the edit distance, or phonetic confusability, of the candidate hotword to other words or phrases in the language of the speaker. To identify the edit distance, for example, finite state transducers, phone-to-phone transducers, and/or language models may be used to determine the most likely words or phrases that may be confused with the candidate hotword. In the present example, it may be determined that several words and phrases such as "visa", "piece of", and "beet saw" (among others) have a high likelihood of being confused with the candidate hotword "pizza". Based on an edit distance criteria, for example, one of the hotword strength evaluation engines 120 may determine that a moderate quantity of words or phrases have a small edit distance from the candidate hotword "pizza", and thus may produce a corresponding feature score 124b (e.g., a medium or low score).

[0022] Further, in the present example, one of the hotword strength evaluation engines 120 may evaluate the transcription 116 and/or data provided by the training examples data store 122 to identify a quantity of pronunciations specified for the candidate hotword. To identify possible pronunciations of the word "pizza", for example, the hotword strength evaluation engine may reference the word in a dictionary or pronunciation guide, based on the transcription 116. As another example, the hotword strength evaluation engine may reference the word "pizza" in the training examples data store 122, and may identify possible variations in how the word has been pronounced by various speakers. In the present example, it may be determined that although only one official American English pronunciation of the word "pizza" is found in the dictionary or pronunciation guide, there is some variation in how the word is pronounced by American speakers. For example, some American speakers may use a native Italian-speaker pronunciation of the word "pizza". Based on a quantity of pronunciations criteria, for example, one of the hotword strength evaluation engines 120 may determine that a low to moderate quantity of pronunciations exist for the candidate hotword "pizza", and thus may produce a corresponding feature score 124c (e.g., a high or medium score).

[0023] During stage (C), a hotword suitability score is generated. Upon generating each of the feature scores 124a, 124b, and 124c, for example, a set of feature scores 126 may be provided to a hotword score generator 128. The hotword score generator 128, for example, may aggregate the scores 124a, 124b, and 124c and provide the aggregated feature scores as a hotword suitability score. To generate the hotword suitability score, for example, the hotword score generator 128 may use logistic regression or an alternative classifier to train a model that estimates a confidence value for the candidate hotword (e.g., the word "pizza"). In general, a high confidence

value may be associated with a potentially suitable hotword. As another example, the hotword score generator may assign appropriate weights (e.g., based on empirical analysis) to each of the feature scores 124a, 124b, and 124c, and may perform an aggregation operation.

[0024] During stage (D), a representation of the hotword suitability score is provided for display to the user. For example, the computing device 102 (shown here as device 102b) may provide one or more indications (e.g., visual and/or audio) of hotword suitability to the user via the interface 104 (shown here as interface 104b). In the present example, the interface 104b may present to the user a message (e.g., the candidate hotword "pizza" is relatively "weak"), and one or more textual or graphical representations (e.g., line graphs, bar graphs, quadrant charts, etc.) of feature and/or aggregate scores associated with the various hotword scoring criteria (e.g., an aggregate hotword suitability score of 12/100). Moreover, in the present example, the user may be presented with a prompt to submit another candidate hotword (e.g., to "try again"). The prompt may be provided, for example, if the hotword suitability score is under a predetermined threshold.

[0025] Referring again to the example data flow, similar to stage (A), during stage (E), the user is again prompted to provide a candidate hotword. In the present example, the user utters the candidate hotword "smorgasbord" and the computing device 102 captures, encodes, and/or generates speech data 130 based on the utterance. Similar to the speech data 110, for example, the speech data 130 may include audio source data 132, a corresponding set of phones 134, and optionally, a transcription 136. Upon capturing, encoding, and/or generating the speech data 130, for example, the data may be provided to one or more hotword strength evaluation engines 120.

[0026] Similar to stage (B), during stage (F), speech data is processed and evaluated. In the present example, the hotword strength evaluation engines 120 may process and evaluate the speech data 130. The hotword strength evaluation engines 120, for example, may evaluate the audio source data 132, the set of phones 134, the transcription 136, and other appropriate speech data associated with the candidate hotword "smorgasbord".

[0027] In the present example, one of the hotword strength evaluation engines 120 may evaluate the audio source data 132 to determine a word length in frames. For example, the candidate hotword "smorgasbord" may be considered to be relatively long compared to other words (e.g., based on an average word length), or the particular speaker may happen to utter the word "pizza" relatively slowly compared to other words or other speakers. Based on a word length criterion, for example, one of the hotword strength evaluation engines 120 may produce a corresponding feature score 124a (e.g., a high score).

[0028] Continuing with the present example, one of the hotword strength evaluation engines 120 may evaluate the set of phones 134 and/or the transcription 136 to iden-

tify the edit distance of the candidate hotword to other words or phrases in the language of the speaker. In the present example, it may be determined that the word "harpichord" has a remote possibility of being confused with the candidate hotword "smorgasbord". Based on an edit distance criteria, for example, one of the hotword strength evaluation engines 120 may determine that a small quantity of words or phrases have a moderate edit distance from the candidate hotword "smorgasbord", and thus may produce a corresponding feature score 124b (e.g., a high score).

[0029] Further, in the present example, one of the hotword strength evaluation engines 120 may evaluate the transcription 136 and/or data provided by the training examples data store 122 to identify a quantity of pronunciations specified for the candidate hotword. Similar to the word "pizza", for example, it may be determined that although only one official American English pronunciation of the word "smorgasbord" is found in a dictionary or pronunciation guide, there is some variation in how the word is pronounced by American speakers. For example, some American speakers may use a Swedish pronunciation of the word "smorgasbord". Based on a quantity of pronunciations criteria, for example, one of the hotword strength evaluation engines 120 may determine that a low to moderate quantity of pronunciations exist for the candidate hotword "smorgasbord", and thus may produce a corresponding feature score 124c (e.g., a high or medium score).

[0030] Similar to stage (C), during stage (G), a hotword suitability score is generated. Upon generating each of the feature scores 124a, 124b, and 124c for the speech data 130, for example, a set of feature scores 126 may be provided to the hotword score generator 128. Further, similar to stage (D), during stage (H), a representation of the hotword suitability score is provided for display to the user. In the present example, the interface 104 (shown here as interface 104c) may present to the user a message (e.g., the candidate hotword "smorgasbord" is relatively "strong"), and one or more textual or graphical representations of feature and/or aggregate scores associated with the various hotword scoring criteria (e.g., an aggregate hotword suitability score of 93/100). Moreover, in the present example, the user may be presented with a message that a candidate hotword is designated as a device hotword (e.g., the hotword is "accepted"). The message may be provided, for example, if the hotword suitability score is over a predetermined threshold.

[0031] As described in the preceding examples, both criteria related to a candidate hotword and to a particular speaker may be considered when evaluating a hotword. The device user may review hotword suitability information provided by the computing device 102 to select a hotword. The use of suitable hotwords may improve hotword detection by the device, thus improving device performance and conserving device resources. Several criteria for evaluating speech data and for generating hotword suitability scores have been described in the pre-

ceding examples. However, as will be described in further examples below, additional criteria, fewer criteria, and/or different criteria may be considered.

[0032] FIG. 2 is a diagram of an example system 200 that may determine the suitability of a spoken hotword. The system 200, for example, may include hardware and/or software components of one or more client computing devices and/or backend servers. In general, the system 200 can receive speech data that encodes a candidate hotword spoken by a user, can evaluate and transcribe the speech data, can generate a hotword suitability score for the candidate hotword, and can provide a representation of the hotword suitability score for display to the user.

[0033] The system 200 may include computer-readable media 202 (e.g., one or more floppy disk devices, hard disk devices, optical disk devices, tape devices, flash memory devices or other similar solid state memory device, or an array of devices) for storing instructions (e.g., application program code) for performing the processes described herein. Further, the system 200 may include one or more processors 204 configured to execute the instructions and to perform various device operations, such as input/output, communication, data processing, and the like.

[0034] The system 200 may include a user interface generator 206 for providing an interface (e.g., graphical, audio, and/or tactile) interface between the system 200 and a user. For example, the user interface generator 206 can provide the interfaces 104a, 104b, and 104c (shown in FIG. 1) for providing prompts, messages, and representations of hotword suitability to the user. Further, the system 200 may include an audio interface 208. The audio interface 208, for example, may include one or more microphones for receiving spoken words and phrases from the user, and one or more speakers for providing audio prompts and feedback (e.g., tones, recordings, computer-generated speech, etc.) to the user. Moreover, the audio interface 208 may provide audio signals associated with received spoken words and phrases to the processors 204 for digital sampling and encoding.

[0035] The system 200 may include an n-gram/phone generator 210. For example, the n-gram/phone generator 210 may analyze encoded speech data provided via the audio interface 208 and the processors 204, and may identify a corresponding series of n-grams and/or phones. The n-grams and/or phones, for example, may be used by a speech recognition engine 212 to perform machine translation and transcription.

[0036] The system 200 may include a plurality of hotword strength evaluation engines 214. Various predetermined criteria may be considered by the hotword strength evaluation engines 214 for evaluating speech data. For example, criteria may include the length of a word or phrase, the quantity of syllables or phones included in the word or phrase, the occurrence of particular phones, the quantity of identified pronunciations for the word or phrase, the variability of word or phrase pronunciation

by the speaker, the availability of training examples for the word or phrase from various speakers, and/or the edit distance between a spoken word or phrase and other words or phrases of a shared language. Each of the hotword strength evaluation engines 214, for example, may evaluate speech data using one or more of the criteria.

[0037] The system 200 may include a score generator 216. For example, a feature score may be determined by the score generator 216 for each of the hotword strength evaluation engines 214, and an overall hotword suitability score may be determined by aggregating the feature scores. A representation of the hotword suitability score may be provided to the user, for example, via an interface generated by the user interface generator 206.

[0038] FIG. 3 is a flow chart illustrating an example process 300 for determining the suitability of a spoken hotword and for providing a representation of the suitability. In some implementations, the process 300 may be performed by the systems 100 and/or 200, and will be described as such for clarity. Briefly, the process 300 includes receiving a spoken hotword, evaluating the hotword, scoring the hotword, and providing a representation of the hotword's suitability.

[0039] In more detail, when the process 300 begins (302), speech data that encodes a candidate hotword spoken by a user is received (304). The user, for example, may want to provide a keyword or key phrase for waking up and/or unlocking his or her stationary device (e.g., a home automation system) or mobile device (e.g., a smartphone). In the present example, the user may choose to provide a phrase that is memorable for the user, yet unlikely to be conceived of by other users (i.e., a voice password). Thus, the user in the present example utters the phrase "pizza smorgasbord".

[0040] Referring to FIG. 2, for example, the audio interface 208 can receive the candidate hotword uttered by the user and can provide audio signals associated with the received utterance to the processors 204 for digital sampling and encoding. Further, the sampled and encoded audio signals can be decomposed into a series of n-grams and phones by the n-gram/phone generator 210, and can be transcribed by the speech recognition engine 212. The speech data and the transcription, for example, may be provided to the hotword strength evaluation engines 214.

[0041] The speech data, or a transcription of the candidate hotword, is evaluated (306), using one or more predetermined criteria. For example, each of the hotword strength evaluation engines 214 may be configured to evaluate the speech data and/or the transcription based on a particular criterion, or a set of related criteria. In general, hotword features or characteristics that may be related to hotword detection performance may be extracted and evaluated.

[0042] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining a word length (e.g., in frames) of the speech data. For example, a framing proc-

ess may be used to divide a speech signal associated with a candidate hotword into shorter speech signals, thereby identifying relatively stable frequency characteristics. In general, longer word lengths (e.g., a larger quantity of frames) may be associated with higher hotword suitability scores. In the present example, the candidate hotword "pizza smorgasbord" may be relatively long in comparison to a typical hotword, and may be associated with a relatively high hotword suitability score when considering a word length feature.

[0043] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining a quantity of syllables or phones associated with the candidate hotword. For example, a quantity of syllables associated with a candidate hotword may be determined by using a transcription of the hotword generated by the speech recognition engine 212 to reference a dictionary including syllable counts for various words. As another example, a quantity of phones associated with a candidate hotword may be determined by evaluating a set of phones identified for the hotword by the n-gram/phone generator 210. In the present example, the candidate hotword "pizza smorgasbord" may be identified as including five syllables in total (e.g., two syllables for the individual word "pizza" and three syllables for the word "smorgasbord"), and may be identified as including ten phones in total (e.g., three phones for the individual word "pizza" and seven phones for the word "smorgasbord"). In general, a larger quantity of syllables or phones may be associated with a higher hotword suitability score. The candidate hotword "pizza smorgasbord" in the present example may have a relatively large quantity of syllables and/or phones in comparison to a typical hotword, and may be associated with a relatively high hotword suitability score when considering a syllable and/or phone quantity feature.

[0044] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining an edit distance between the transcription of the candidate hotword and one or more other words that share a same language as the candidate hotword. Words or phrases having a small edit distance between each other, for example, may generally be confusable with each other, whereas words or phrases having a large edit distance between each other may generally be more easily distinguished. In the present example, it may be determined that the phrase "piece of harpsichord" may have a moderate edit distance from (i.e., may be somewhat confusable with) the candidate hotword "pizza smorgasbord". In general, smaller edit distances may be associated with lower hotword suitability scores. Thus, in the present example, as the phrases "piece of harpsichord" and "pizza smorgasbord" have a moderate edit distance between them, the candidate hotword may be associated with a moderate hotword suitability score when considering an edit distance feature.

[0045] Upon identifying one or more words or phrases

that are within a predetermined edit distance range of the candidate hotword, in some implementations the quantity of such words may be considered when determining an edit distance feature score. For example, the phrase "piece of harpsichord" may have been identified as being within a predetermined edit distance range from the candidate hotword "pizza smorgasbord", but only a few (if any) other phrases may have also been identified as being within the edit distance range. Thus, in the present example, the candidate hotword may be associated with a higher feature score. Further, in some implementations, the likelihoods of a candidate hotword and/or each of the words or phrases that are within the predetermined edit distance range of the candidate hotword (i.e., potentially confusable words or phrases) being used in general conversation may be considered when determining a hotword suitability score. For example, or more of the candidate evaluation engines 214 may determine that the likelihood of the candidate hotword "pizza smorgasbord" and/or the potentially confusable phrase "piece of harpsichord" being spoken in general conversation (e.g., by the device user, and/or by people in general) is relatively low, and the candidate hotword may be associated with an even feature higher score.

[0046] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining a quantity of pronunciations specified for the candidate hotword. For example, a quantity of pronunciations associated with a candidate hotword may be determined by using a transcription of the hotword generated by the speech recognition engine 212 to reference a dictionary including common or accepted pronunciations of the hotword. As another example, recorded instances of the user and/or other users having spoken the candidate hotword may be evaluated. In general, fewer quantities of pronunciations may be associated with higher hotword suitability scores. For example, one of the hotword strength evaluation engines 214 may identify a single official pronunciation of the phrase "pizza smorgasbord" from a dictionary source, but may identify a small to medium quantity of different pronunciations of the phrase by a population of speakers. Thus, in the present example, the candidate hotword "pizza smorgasbord" may be associated with a relatively high to medium hotword suitability score when considering a quantity of pronunciations feature.

[0047] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining that one or more particular phones or syllables occur in the transcription of the candidate hotword. To identify the particular syllables, for example, a transcription of the candidate hotword may be used to reference a dictionary or pronunciation guide, and the defined syllables associated with the hotword may be identified. As another example, one of the hotword strength evaluation engines 214 may evaluate phones identified for a spoken candidate hotword by the n-gram/phone generator 210. In general, the occurrence

of certain phones or syllables in the transcription may be associated with a higher hotword suitability score. For example, the second syllable of the word "pizza" (i.e., "suh") included in the phrase "pizza smorgasbord" may be associated with a high detectability confidence level, and a relatively high hotword suitability. However, the occurrence of certain other phones in the transcription may generally be associated with a lower hotword suitability score. For example, the first syllable of the word "pizza" (i.e., "peet") may be associated with a low detectability confidence level (e.g., it may be confusable with the syllable "beet"), and a relatively low hotword suitability. Upon determining and evaluating the particular phones or syllables that occur in a candidate hotword, for example, the constituent scores associated with each phone or syllable may be used for determining an aggregated score. For example, the candidate hotword "pizza smorgasbord" may include particular phones or syllables associated with low, medium, and high scores. Thus, the candidate hotword in the present example may be associated with a medium hotword suitability score when considering particular phone and/or syllable features.

[0048] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining a quantity of training examples of the hotword that are available to a speech recognizer. Referring to FIG. 1, for example, recorded instances of various individuals (e.g., the device user and/or other users) having spoken various words and phrases may be maintained by the training examples data store 122. Moreover, the training examples data store 122 may include transcription and/or index data, for example, so recorded instances of particular words and/or phrases may be referenced and analyzed. In general, a smaller quantity of training examples may be associated with a lower hotword suitability score. The candidate hotword "pizza smorgasbord" in the present example may have a relatively small quantity of training examples in comparison to a typical hotword, and may be associated with a relatively low hotword suitability score when considering a quantity of training examples feature.

[0049] If no training examples are found for a particular candidate hotword, in some implementations a nearest neighbor hotword (e.g., a word or phrase associated with similar phones, syllables, or other appropriate features) may be evaluated. For example, a user may choose to provide an invented word (i.e., a word that does not exist in a dictionary) as a candidate hotword. In the present example, as training examples may be unavailable for the invented candidate hotword, a training examples feature score for the hotword may be low, but other feature scores may be relatively high, thus offsetting the low feature score when determining an aggregate hotword suitability score. For example, a particular invented word may include one or more syllables or phones that are associated with high feature scores, and the invented word may be associated with a high score when considering an edit distance feature.

[0050] In some implementations, evaluating the speech data and/or a transcription of the candidate hotword may include determining an extent to which pronunciation of the candidate hotword varies among multiple utterances of the candidate hotword by the user. Referring to FIG. 1, for example, the user may be prompted by the computing device 102a to "please say your desired hotword" multiple times, and the various instances of the spoken hotword (e.g., the associated audio signals) may be compared with each other for consistency by one or more of the hotword strength evaluation engines 120. In the present example, the user may be somewhat inconsistent when pronouncing the candidate hotword "pizza smorgasbord" - the phrase may be difficult for the particular speaker to pronounce consistently. In general, a greater pronunciation variation extent may be associated with a lower hotword suitability score. Thus, in the present example, the candidate hotword "pizza smorgasbord" may be associated with a relatively low hotword suitability score when considering a pronunciation variation feature.

[0051] Various possible hotword features, criteria for evaluating speech data, and techniques for generating hotword suitability scores have been described in the preceding examples. However, additional, fewer, and/or different features, criteria, and/or techniques may be employed. In general, classifiers may be trained to identify and consider potentially influential features and criteria related to speech data. Moreover, techniques for generating hotword suitability scores may change over time as additional speech data (e.g., training examples) are considered.

[0052] A hotword suitability score for the candidate hotword is generated (308), based on evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria (e.g., the criteria described in the preceding examples). The hotword suitability score may reflect a suitability of the candidate hotword for use as a hotword. For example, a high hotword suitability score may be associated with a candidate hotword that is likely to perform well when used to activate a computing device, whereas a low hotword suitability score may be associated with a candidate hotword that is likely to perform poorly. Well performing hotwords, for example, may be readily detectable and unlikely to be confused with natural language conversation. Poorly performing hotwords, for example, may be difficult for a computing device to detect and may be likely to be confused with other words and phrases.

[0053] In some implementations, generating a hotword suitability score may include generating, for each of two or more predetermined criteria, a feature score based on an evaluation of the speech data (or of the transcription of the candidate hotword) using the predetermined criteria, aggregating the feature scores, and outputting, as the hotword suitability score, the aggregated feature scores. In the present example, the candidate hotword "pizza smorgasbord" may be associated with high or me-

dium feature scores for a word length criterion, a quantity of syllables or phones criterion, an edit distance criterion, a quantity of pronunciations criterion, and an occurrence of particular phones or syllables criterion. However, the candidate hotword "pizza smorgasbord" in the present example may also be associated with low feature scores for a quantity of training examples criterion and for a pronunciation variation criterion. An aggregate hotword suitability score, for example, may depend on how weights may be applied to the various feature scores, and/or may depend on how classifiers for generating the scores are trained. In the present example, the candidate hotword "pizza smorgasbord" may be associated with a medium to high hotword suitability score.

[0054] A representation of the hotword suitability score is provided (310) for display to the user, thereby ending the process (312). Referring to FIG. 1, for example, the interfaces 104b and 104c are shown for displaying representations of hotword suitability scores to users. The interface 104b, for example, shows a representation of a hotword suitability score where the representation includes data indicating that the candidate hotword is unsuitable as a hotword, and that the candidate hotword has been rejected. The interface 104c, for example, shows a representation of a hotword suitability score where the representation includes data indicating that the candidate hotword is suitable as a hotword, and that the candidate hotword has been accepted. However, in some implementations, hotword suitability scores may serve as general information about how well hotwords are predicted to perform on a computing device, and a user may choose to act upon or to ignore the information.

[0055] Various representations (e.g., visual and/or audio) of a hotword suitability score may be provided to the user. In some implementations, a text score (e.g., a ten point scale, a hundred point scale, a letter grade, etc.) may be provided. For example, an interface may display the letter grade "B+" for the candidate hotword "pizza smorgasbord". In some implementations, a representation of a hotword suitability score may include a color-coded representation of the score. Similar to a traffic light, for example, the color green may be associated with high suitability scores, the color yellow may be associated with medium suitability scores, and the color red may be associated with low suitability scores. In some implementations, a representation of a hotword suitability score may include a bar, a gauge, or a meter. A high overall suitability score, for example, may be represented with a proportionally high (or long) bar indicator. As another example, each of the individual feature scores included in an aggregate hotword suitability score may be represented by a separate bar, gauge, or meter. Thus, in the present example, if the user is informed that the candidate hotword "pizza smorgasbord" has received relatively high scores for most features, with the exception of a training example criterion and a pronunciation variation criterion, he or she may use the information to improve the candidate hotword (e.g., by providing more training

examples and/or a clearer pronunciation), or to select a different candidate hotword.

[0056] FIG. 4 shows an example of a generic computer device 400 and a generic mobile computer device 450, which may be used with the techniques described here. Computing device 400 is intended to represent various forms of digital computers, such as laptops, desktops, workstations, personal digital assistants, servers, blade servers, mainframes, and other appropriate computers. Computing device 450 is intended to represent various forms of mobile devices, such as personal digital assistants, cellular telephones, smartphones, and other similar computing devices. The components shown here, their connections and relationships, and their functions, are meant to be exemplary only, and are not meant to limit implementations of the inventions described and/or claimed in this document.

[0057] Computing device 400 includes a processor 402, memory 404, a storage device 406, a high-speed interface 408 connecting to memory 404 and high-speed expansion ports 410, and a low speed interface 412 connecting to low speed bus 414 and storage device 406. Each of the components 402, 404, 406, 408, 410, and 412, are interconnected using various busses, and may be mounted on a common motherboard or in other manners as appropriate. The processor 402 may process instructions for execution within the computing device 400, including instructions stored in the memory 404 or on the storage device 406 to display graphical information for a GUI on an external input/output device, such as display 416 coupled to high speed interface 408. In other implementations, multiple processors and/or multiple buses may be used, as appropriate, along with multiple memories and types of memory. Also, multiple computing devices 400 may be connected, with each device providing portions of the necessary operations (e.g., as a server bank, a group of blade servers, or a multi-processor system).

[0058] The memory 404 stores information within the computing device 400. In one implementation, the memory 404 is a volatile memory unit or units. In another implementation, the memory 404 is a non-volatile memory unit or units. The memory 404 may also be another form of computer-readable medium, such as a magnetic or optical disk.

[0059] The storage device 406 is capable of providing mass storage for the computing device 400. In one implementation, the storage device 406 may be or contain a computer-readable medium, such as a floppy disk device, a hard disk device, an optical disk device, or a tape device, a flash memory or other similar solid state memory device, or an array of devices, including devices in a storage area network or other configurations. A computer program product may be tangibly embodied in an information carrier. The computer program product may also contain instructions that, when executed, perform one or more methods, such as those described above. The information carrier is a computer- or machine-readable me-

dium, such as the memory 404, the storage device 406, memory on processor 402, or a propagated signal.

[0060] The high speed controller 408 manages bandwidth-intensive operations for the computing device 400, while the low speed controller 412 manages lower bandwidth-intensive operations. Such allocation of functions is exemplary only. In one implementation, the high-speed controller 408 is coupled to memory 404, display 416 (e.g., through a graphics processor or accelerator), and to high-speed expansion ports 410, which may accept various expansion cards (not shown). In the implementation, low-speed controller 412 is coupled to storage device 406 and low-speed expansion port 414. The low-speed expansion port, which may include various communication ports (e.g., USB, Bluetooth, Ethernet, wireless Ethernet) may be coupled to one or more input/output devices, such as a keyboard, a pointing device, a scanner, or a networking device such as a switch or router, e.g., through a network adapter.

[0061] The computing device 400 may be implemented in a number of different forms, as shown in the figure. For example, it may be implemented as a standard server 420, or multiple times in a group of such servers. It may also be implemented as part of a rack server system 424. In addition, it may be implemented in a personal computer such as a laptop computer 422. Alternatively, components from computing device 400 may be combined with other components in a mobile device (not shown), such as device 450. Each of such devices may contain one or more of computing device 400, 450, and an entire system may be made up of multiple computing devices 400, 450 communicating with each other.

[0062] Computing device 450 includes a processor 452, memory 464, an input/output device such as a display 454, a communication interface 466, and a transceiver 468, among other components. The device 450 may also be provided with a storage device, such as a microdrive or other device, to provide additional storage. Each of the components 450, 452, 464, 454, 466, and 468, are interconnected using various buses, and several of the components may be mounted on a common motherboard or in other manners as appropriate.

[0063] The processor 452 may execute instructions within the computing device 450, including instructions stored in the memory 464. The processor may be implemented as a chipset of chips that include separate and multiple analog and digital processors. The processor may provide, for example, for coordination of the other components of the device 450, such as control of user interfaces, applications run by device 450, and wireless communication by device 450.

[0064] Processor 452 may communicate with a user through control interface 458 and display interface 456 coupled to a display 454. The display 454 may be, for example, a TFT LCD (Thin-Film-Transistor Liquid Crystal Display) or an OLED (Organic Light Emitting Diode) display, or other appropriate display technology. The display interface 456 may comprise appropriate circuitry for driv-

ing the display 454 to present graphical and other information to a user. The control interface 458 may receive commands from a user and convert them for submission to the processor 452. In addition, an external interface 462 may be provide in communication with processor 452, so as to enable near area communication of device 450 with other devices. External interface 462 may provide, for example, for wired communication in some implementations, or for wireless communication in other implementations, and multiple interfaces may also be used.

[0065] The memory 464 stores information within the computing device 450. The memory 464 may be implemented as one or more of a computer-readable medium or media, a volatile memory unit or units, or a non-volatile memory unit or units. Expansion memory 474 may also be provided and connected to device 450 through expansion interface 472, which may include, for example, a SIMM (Single In Line Memory Module) card interface. Such expansion memory 474 may provide extra storage space for device 450, or may also store applications or other information for device 450. Specifically, expansion memory 474 may include instructions to carry out or supplement the processes described above, and may include secure information also. Thus, for example, expansion memory 474 may be provide as a security module for device 450, and may be programmed with instructions that permit secure use of device 450. In addition, secure applications may be provided via the SIMM cards, along with additional information, such as placing identifying information on the SIMM card in a non-hackable manner.

[0066] The memory may include, for example, flash memory and/or NVRAM memory, as discussed below. In one implementation, a computer program product is tangibly embodied in an information carrier. The computer program product contains instructions that, when executed, perform one or more methods, such as those described above. The information carrier is a computer- or machine-readable medium, such as the memory 464, expansion memory 474, memory on processor 452, or a propagated signal that may be received, for example, over transceiver 468 or external interface 462.

[0067] Device 450 may communicate wirelessly through communication interface 466, which may include digital signal processing circuitry where necessary. Communication interface 466 may provide for communications under various modes or protocols, such as GSM voice calls, SMS, EMS, or MMS messaging, CDMA, TDMA, PDC, WCDMA, CDMA2000, or GPRS, among others. Such communication may occur, for example, through radio-frequency transceiver 468. In addition, short-range communication may occur, such as using a Bluetooth, WiFi, or other such transceiver (not shown). In addition, GPS (Global Positioning System) receiver module 470 may provide additional navigation- and location-related wireless data to device 450, which may be used as appropriate by applications running on device 450.

[0068] Device 450 may also communicate audibly using audio codec 460, which may receive spoken information from a user and convert it to usable digital information. Audio codec 460 may likewise generate audible sound for a user, such as through a speaker, e.g., in a handset of device 450. Such sound may include sound from voice telephone calls, may include recorded sound (e.g., voice messages, music files, etc.) and may also include sound generated by applications operating on device 450.

[0069] The computing device 450 may be implemented in a number of different forms, as shown in the figure. For example, it may be implemented as a cellular telephone 480. It may also be implemented as part of a smart-phone 482, personal digital assistant, or other similar mobile device.

[0070] Various implementations of the systems and techniques described here may be realized in digital electronic circuitry, integrated circuitry, specially designed ASICs (application specific integrated circuits), computer hardware, firmware, software, and/or combinations thereof. These various implementations may include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which may be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device.

[0071] These computer programs (also known as programs, software, software applications or code) include machine instructions for a programmable processor, and may be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the terms "machine-readable medium" "computer-readable medium" refers to any computer program product, apparatus and/or device (e.g., magnetic discs, optical disks, memory, Programmable Logic Devices (PLDs)) used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine instructions as a machine-readable signal. The term "machine-readable signal" refers to any signal used to provide machine instructions and/or data to a programmable processor.

[0072] To provide for interaction with a user, the systems and techniques described here may be implemented on a computer having a display device (e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor) for displaying information to the user and a keyboard and a pointing device (e.g., a mouse or a trackball) by which the user may provide input to the computer. Other kinds of devices may be used to provide for interaction with a user as well; for example, feedback provided to the user may be any form of sensory feedback (e.g., visual feedback, auditory feedback, or tactile feedback); and input from the user may be received in any form, including acoustic, speech, or tactile input.

[0073] The systems and techniques described here may be implemented in a computing system that includes a back end component (e.g., as a data server), or that includes a middleware component (e.g., an application server), or that includes a front end component (e.g., a client computer having a graphical user interface or a Web browser through which a user may interact with an implementation of the systems and techniques described here), or any combination of such back end, middleware, or front end components. The components of the system may be interconnected by any form or medium of digital data communication (e.g., a communication network). Examples of communication networks include a local area network ("LAN"), a wide area network ("WAN"), and the Internet.

[0074] The computing system may include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[0075] A number of embodiments have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. In addition, the logic flows depicted in the figures do not require the particular order shown, or sequential order, to achieve desirable results. In addition, other steps may be provided, or steps may be eliminated, from the described flows, and other components may be added to, or removed from, the described systems. Accordingly, other embodiments are within the scope of the following claims.

[0076] This specification also includes the following clauses:

1. A computer-implemented method comprising:

receiving speech data that encodes a candidate hotword spoken by a user;
evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria;
generating a hotword suitability score for the candidate hotword based on evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria; and
providing a representation of the hotword suitability score for display to the user.

2. The method of clause 1, wherein the hotword suitability score reflects a suitability of the candidate hotword for use as a hotword.

3. The method of clause 2, wherein a hotword is a keyword or a key phrase for waking up a device.

4. The method of clause 2, wherein a hotword is a keyword or a key phrase for unlocking a device.

5. The method of clause 2, wherein a hotword is a voice password.

6. The method of clause 1, wherein the speech data is received from a mobile device that is operated by the user.

7. The method of clause 1, wherein generating a hotword suitability score comprises:

generating, for each of two or more predetermined criteria, a feature score based on an evaluation of the speech data, or of the transcription of the candidate hotword, using the predetermined criteria;
aggregating the feature scores; and
outputting, as the hotword suitability score, the aggregated feature scores.

8. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining a word length, in frames, of the speech data, wherein longer word lengths are associated with higher hotword suitability scores.

9. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining a quantity of syllables or phones associated with the candidate hotword, wherein a larger quantity of syllables or phones is associated with a higher hotword suitability score.

10. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining an edit distance between the transcription of the candidate hotword and one or more other words that share a same language as the candidate hotword, wherein smaller edit distances are associated with lower hotword suitability scores.

11. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining a quantity of pronunciations specified for the candidate hotword,

wherein fewer quantities of pronunciations are associated with higher hotword suitability scores.

12. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining that one or more particular phones occur in the transcription of the candidate hotword.

13. The method of clause 12, wherein the occurrence of the particular phones in the transcription is associated with a higher hotword suitability score.

14. The method of clause 12, wherein the occurrence of the particular phones in the transcription is associated with a lower hotword suitability score.

15. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining a quantity of training examples of the hotword that are available to a speech recognizer,
wherein a smaller quantity of training examples is associated with a lower hotword suitability score.

16. The method of clause 1, wherein evaluating the speech data, or a transcription of the candidate hotword, comprises:

determining an extent to which pronunciation of the candidate hotword varies among multiple utterances of the candidate hotword by the user, wherein a greater extent is associated with a lower hotword suitability score.

17. The method of clause 1, wherein the representation of the hotword suitability score comprises data indicating that the candidate hotword is suitable as a hotword.

18. The method of clause 1, wherein the representation of the hotword suitability score comprises data indicating that the candidate hotword has been accepted as a hotword.

19. The method of clause 1, wherein the representation of the hotword suitability score comprises a color-coded representation of the hotword suitability score.

20. The method of clause 1, wherein the representation of the hotword suitability score comprises a

bar, a gauge or a meter.

21. A system comprising:

one or more computers and one or more storage devices storing instructions that are operable, when executed by the one or more computers, to cause the one or more computers to perform operations comprising:

receiving speech data that encodes a candidate hotword spoken by a user; evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria; generating a hotword suitability score for the candidate hotword based on evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria; and providing a representation of the hotword suitability score for display to the user.

22. A computer-readable storage device storing software comprising instructions executable by one or more computers which, upon such execution, cause the one or more computers to perform operations comprising:

receiving speech data that encodes a candidate hotword spoken by a user; evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria; generating a hotword suitability score for the candidate hotword based on evaluating the speech data, or a transcription of the candidate hotword, using one or more predetermined criteria; and providing a representation of the hotword suitability score for display to the user.

Claims

1. A computer-implemented method comprising:

receiving speech data indicative of a candidate hotword; determining that the candidate hotword satisfies one or more hotword suitability criteria; and in response to determining that the candidate hotword satisfies one or more hotword suitability criteria, determining whether subsequently received audio data includes an utterance of the hotword.

2. The computer-implemented method of claim 1, fur-

ther comprising:

determining that the candidate hotword includes a particular quantity of syllables or phones; and wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based at least on the particular quantity of syllables or phones, that the candidate hotword satisfies one or more hotword suitability criteria.

3. The computer-implemented method of claim 1 or 2, further comprising:

determining that the candidate hotword includes one or more particular phones; and wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based at least on the one or more particular phones included in the candidate hotword, that the candidate hotword satisfies one or more hotword suitability criteria.

4. The computer-implemented method of any preceding claim, further comprising:

determining that the candidate hotword is a particular edit distance from one or more terms; and wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based at least on the particular edit distance between the candidate hotword and the one or more terms, that the candidate hotword satisfies one or more hotword suitability criteria.

5. The computer-implemented method of any preceding claim, further comprising:

determining that the candidate hotword corresponds to a particular quantity of pronunciations; and wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based at least on the particular quantity of pronunciations, that the candidate hotword satisfies one or more hotword suitability criteria.

6. The computer-implemented method of any preceding claim, further comprising:

before determining that the candidate hotword satisfies one or more hotword suitability criteria, receiving audio data that encodes an utterance of one or more words, at least a portion of which represent the candidate hotword.

7. The computer-implemented method of claim 6, further comprising:

determining that the candidate hotword is encoded in a particular quantity of frames of the audio data that encodes the utterance of one or more words; and
 wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based at least on the particular quantity of frames, that the candidate hotword satisfies one or more hotword suitability criteria.

8. The computer-implemented method of claim 6, wherein receiving audio data that encodes the utterance of one or more words comprises:

receiving audio data that encodes an utterance of one or more words spoken by a user of a client device.

9. The computer-implemented method of claim 8, further comprising:

in response to determining that subsequently received audio data includes an utterance of the hotword, providing, to an application running on the client device, one or more commands for waking up the client device.

10. The computer-implemented method of claim 8, further comprising:

in response to determining that subsequently received audio data includes an utterance of the hotword, providing, to an application running on the client device, one or more commands for unlocking the client device.

11. The computer-implemented method of claim 8, further comprising:

in response to determining that subsequently received audio data includes an utterance of the hotword, providing, to an application running on the client device, one or more commands for applying settings of the user's preference.

12. The computer-implemented method of claim 8, further comprising:

in response to determining that the candidate hotword satisfies one or more hotword suitability criteria, providing an indication that the candidate hotword satisfies one or more hotword suitability criteria for output at the client device.

13. The computer-implemented method of claim 8, further comprising:

obtaining additional audio data encoding one or more utterances of the candidate hotword spoken by the user of the client device;
 determining, based on the received audio data and the additional audio data, that the user's pronunciation of the candidate hotword varies to a particular extent from utterance to utterance; and
 wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based at least on the particular extent to which the user's pronunciation of the candidate hotword varies from utterance to utterance, that the candidate hotword satisfies one or more hotword suitability criteria.

14. The computer-implemented method of any preceding claim, further comprising:

estimating a confidence value that reflects a level of confidence in the candidate hotword representing a suitable hotword; and
 wherein determining that the candidate hotword satisfies one or more hotword suitability criteria, comprises determining, based on the confidence value, that the candidate hotword satisfies one or more hotword suitability criteria.

15. The computer-implemented method of claim 14, further comprising:

generating one or more textual or graphical representations of the confidence value that reflects a level of confidence in the candidate hotword representing a suitable hotword; and
 providing the one or more textual or graphical representations for output at one or more computing devices.

16. The computer-implemented method of claim 15, further comprising:

generating a textual or graphical indication that the candidate hotword satisfies one or more hotword suitability criteria; and
 providing the textual or graphical indication for output at the one or more computing devices.

17. The computer-implemented method of any preceding claim, further comprising:

determining that a particular quantity of training examples of the candidate hotword are available to the speech recognizer; and
 wherein determining that the candidate hotword

satisfies one or more hotword suitability criteria, comprises determining, based at least on the particular quantity of training examples, that the candidate hotword satisfies one or more hotword suitability criteria.

5

18. A computer program comprising computer-executable instructions which, when executed by at least one computing apparatus, cause the method of any preceding claim to be performed.

10

19. A computer program product, encoded on one or more non-transitory computer storage media, comprising instructions that when executed by one or more computers cause the one or more computers to perform the method of any of claims 1 to 17.

15

20. A system comprising one or more computers configured to perform the method of any of claims 1 to 17.

20

25

30


35


40

45

50

55





200

FIG. 2

300 ↗




FIG. 3




FIG. 4



EUROPEAN SEARCH REPORT

 Application Number
 EP 17 19 8918

5

10

15

20

25

30

35

40

45

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 5 754 977 A (GARDNER MARC A [US] ET AL) 19 May 1998 (1998-05-19)	1,4,6, 8-10,12, 14-16, 18-20	INV. G10L17/24 G06F21/46 G10L15/08
A	* column 1, line 53 - column 2, line 29 * * column 3, line 14 - column 4, line 47 * * column 5, line 29 - line 31; figures 3,4 *	2,3,5,7, 11,13,17	G10L15/22 G10L25/51 G06F21/32
X,P	----- A. ZGANK ET AL: "Predicting the Acoustic Confusability between Words for a Speech Recognition System using Levenshtein Distance", ELECTRONICS AND ELECTRICAL ENGINEERING, vol. 18, no. 8, 18 October 2012 (2012-10-18), XP055082737, ISSN: 1392-1215, DOI: 10.5755/j01.eee.18.8.2628	1,4,6, 8-10,12, 14-16, 18-20	
A,P	* page 81, left-hand column, line 7 - line 20 * * page 81, right-hand column, paragraph II - page 82, left-hand column * * page 83, left-hand column, paragraph V. - right-hand column *	2,3,5,7, 11,13,17	TECHNICAL FIELDS SEARCHED (IPC) G10L G06F
----- -/--			
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 24 January 2018	Examiner Ebbinghaus, Stefanie
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons ----- & : member of the same patent family, corresponding document	

 2
 EPO FORM 1503 03.82 (P04C01)



EUROPEAN SEARCH REPORT

 Application Number
EP 17 19 8918

5

10

15

20

25

30

35

40

45

2

50

55

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	Jan Anguity ET AL: "Word Confusability Prediction in Automatic Speech Recognition", Proceedings of nternational Conference on Spoken Language Processing (Interspeech), 4 October 2004 (2004-10-04), pages 1489-1492, XP055082735, Jeju Island, Korea Retrieved from the Internet: URL:http://nlp.lsi.upc.edu/papers/anguita04.pdf [retrieved on 2013-10-07]	1,4,6,8-10,12,14-16,18-20	TECHNICAL FIELDS SEARCHED (IPC)
A	* page 1, right-hand column, paragraph 2.1 - page 2, left-hand column * * page 2, right-hand column, paragraph 4.1 - page 3, right-hand column, paragraph 4.2 *	2,3,5,7,11,13,17	
X	ROE D B ET AL: "PREDICTION OF WORD CONFUSABILITIES FOR SPEECH RECOGNITION", ICSLP 94 : 1994 INTERNATIONAL CONFERENCE ON SPOKEN LANGUAGE PROCESSING. YOKOHAMA, JAPAN, SEPT. 18 - 22, 1994; [INTERNATIONAL CONFERENCE ON SPOKEN LANGUAGE PROCESSING. (ICSLP)], YOKOHAMA : ASJ, JP, vol. 1, 18 September 1994 (1994-09-18), pages 227-230, XP000855228,	1,4,6,8-10,12,14-16,18-20	
A	* page 227, right-hand column, line 6 - page 228, left-hand column, line 21 *	2,3,5,7,11,13,17	
A	US 2004/250139 A1 (HURLEY JOHN C [US]) 9 December 2004 (2004-12-09) * figures 2-6 *	1-20	
A	WO 2008/148659 A1 (IBM [US]; IBM UK [GB]; KEOHANE SUSANN MARIE [US]; MCBREARTY GERALD FRA) 11 December 2008 (2008-12-11) * figures 3A, 3B *	1-20	
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 24 January 2018	Examiner Ebbinghaus, Stefanie
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 17 19 8918

5

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

24-01-2018

10

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5754977 A	19-05-1998	NONE	
US 2004250139 A1	09-12-2004	US 2004250139 A1	09-12-2004
		US 2008072320 A1	20-03-2008
WO 2008148659 A1	11-12-2008	BR PI0811365 A2	16-06-2015
		CN 101681412 A	24-03-2010
		JP 5337149 B2	06-11-2013
		JP 2010529544 A	26-08-2010
		KR 20100028522 A	12-03-2010
		TW 200905625 A	01-02-2009
		US 2008307235 A1	11-12-2008
		WO 2008148659 A1	11-12-2008

15

20

25

30

35

40

45

50

55

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 61667817 A [0001]
- US 56757212 A [0001]